

---

## **Rule WLM222:      Service Class was Active, but server was CPU capped**

---

**Finding:** CPExpert has determined that resource capping was a major cause of the service class not achieving its performance goal.

**Impact:** The impact of this finding depends upon the amount of resource capping delay experienced by the service class. A high percent of resource capping delay means HIGH IMPACT while a low percent of resource capping means LOW IMPACT. See the output associated with the rule which caused this rule to be invoked (Rule WLM104 or Rule WLM105, depending upon the type of service class and performance goal).

**Logic flow:** The following rules cause this rule to be invoked:  
Rule WLM120: Significant transaction time was in Active state  
Rule WLM121: Significant transaction time was in Ready state

**Discussion:** When CPExpert produces Rule WLM104 or Rule WLM105 to indicate that a subsystem service class did not achieve its performance goal, the logic of these rules tries to identify the cause of the delay. The cause of the delay initially is analyzed from the "served" service class view. Rule WLM120(series) to Rule WLM130(series) describe the results from this analysis.

After analyzing the subsystem transaction delays, CPExpert identifies the service classes which serve the transactions. The subsystem transactions typically are CICS transactions, and the servers are the CICS regions. Alternatively, the transactions could be IMS transactions and the servers could be the IMS control regions or transaction processing regions.

Address spaces executing in the system can be in a variety of states from the perspective of the Workload Manager: using the CPU, delayed for an identifiable reason, or delayed for some unknown reason.

The System Resources Manager (SRM) periodically samples the state of each address space in each service class. These samples are accumulated into variables which are recorded by RMF in the "Service Class Period Data Section" of SMF Type 72 (Subtype 3) records. Please see Section 4 for a discussion of these states and the sampling process.

CPExpert produces Rule WLM120 when a significant cause of delay to a subsystem transaction was that the transaction was in Active state. The

---

Active state indicates that a task was executing on behalf of the transaction, from the perspective of CICS or IMS.

CPEXpert produces Rule WLM121 when a significant cause of delay to a subsystem transaction was that the transaction was in Ready state. The ready state indicates that there was a program ready to execute on behalf of a work request in the "served" service class, but that the work manager has given priority to another work request. In the case of a CICS region, this means that there were more CICS tasks ready to execute in the "served" service class than were dispatched by CICS.

When Rule WLM120 or WLM121 are produced, CPEXpert analyzes the CPU requirements and CPU capping of the server service class to determine whether the transaction required a significant amount of CPU time (which might be indicated by the Active state) or whether transactions were delayed in the Ready state (because the CICS region had been CPU capped).

CPU usage and other resource requirements are not contained in the SMF Type 72 records which describe the subsystem transaction service class. These subsystem transaction service classes are not address spaces, but are logical groupings of transactions. Resource information is not recorded by SMF for the transactions, but is recorded for the address spaces (the servers) providing service to the service classes. Consequently, CPEXpert analyzes the resource requirements of the server service classes.

CPEXpert analyzes the amount of CPU capping time of the server service classes by the following process:

- CPEXpert first computes the number of samples which found an address space executing in the service class. This is done by summing CPU Using samples (R723CCUS), Total Wait samples (R723CTOT), and Unknown Delay samples (R723CUNK). The result is titled "EXSAMP" in the code.
- CPEXpert divides the number of CPU Capping samples (R723CCCA) by the EXSAMP value, to yield the percent of execution samples in which the SRM found an address space was CPU capped. The average transaction response time is multiplied by the resulting percentage to yield the amount of time when the average transaction was delayed because of CPU capping.
- Server service classes might serve multiple subsystem service classes. For example, a CICS region (the server) might provide service to a number of service classes which describe different CICS transactions. If the server provides service to multiple service classes, the above

---

technique automatically "pro-rates" the CPU capping delay of the server. This automatic "pro-rating" is possible because the sampling process of the SRM is independent of the transactions executing.

Resource capping is a way of controlling the distribution of CPU service to one or more service classes. Resource capping is implemented by defining "resource groups" to the Workload Manager. A resource group is simply a named set of two values: a minimum CPU service specification and a maximum CPU service specification. The specifications are in terms of **unweighted CPU service units** (that is, the CPU service coefficients are not applied to TCB nor SRB raw CPU service units).

The Workload Manager will attempt to provide the minimum CPU service to the resource group and will restrict the resource from using more than the maximum CPU service.

Service classes are associated with resource groups; however, a particular service class can be associated with only one resource group<sup>1</sup>.

**It normally is not advisable to use resource groups.** IBM provides the facility solely for special cases, and IBM does not contemplate resource groups being normally used.

Resource group specifications are "preemptive" in nature, in that the Workload Manager attempts to honor resource group specifications before considering other service specifications. Consequently, **resource group specifications could nullify the rest of the Workload Manager's algorithms.**

When the maximum CPU service specified in the resource group has been used, the Workload Manager marks "non-dispatchable" the TCBs and SRBs associated with the service classes assigned to the resource group. This is the situation addressed by Rule WLM222.

As the System Resources Manager takes its samples of the state of address spaces, it examines whether a dispatchable unit (TCB or SRB) is marked non-dispatchable because of a resource group maximum. Samples reflecting the resource group maximum are recorded by RMF in the SMF Type 72 delay samples, as CPU Capping Delay (R723CCCA).

As described earlier, CPExpert computes the percent of CPU Capping Delay for the server service class, as a function of the overall execution of transactions served by the server service class. CPExpert produces Rule

---

<sup>1</sup>Please see Section 4 (Chapter 1.6) for a discussion of resource groups and how the Workload Manager implements the resource group specifications.

---

WLM222 if the percent of CPU Capping Delay for the server service class is greater than the significance value specified in the **WLM SIG** guidance variable in USOURCE(WLMGUIDE).

With Rule WLM222, CPExpert provides the total number of ending transactions in the RMF measurement interval, the total CPU service units required to service the transactions, the average CPU service units per transaction, and the average percent resource capping delay to transactions active in the service class.

**Suggestion:** As mentioned above, resource groups are intended for very special situations. In most environments, it is far better to allow the Workload Manager to manage system resources to meet the performance goals specified for various service classes. Using resource groups takes control away from the Workload Manager.

Further, specifying maximum CPU service units may result in unused CPU capacity if there are no other service classes ready to use the CPU service.

CPExpert suggests that you consider the following alternatives:

- While there may be unusual situations in which control must be removed from the Workload Manager, please consider whether you have such an unusual situation. If you do not have an unusual situation, you may wish to remove the resource group from the service class. This is particularly true since the service class missing its performance goal describes response goals.
- Alternatively, you should review the performance goal specified for the service class identified by Rule WLM222. CPExpert performs "delay analysis" only on service classes which fail to achieve their performance goal. Consequently, the service class identified by Rule WLM222 had failed to achieve its performance goal.

The performance goal may be incompatible with the resource group Capacity Maximum, and you may wish to either increase the performance goal (for response goals) or decrease the performance goal (for execution velocity goals).

- Alternatively, you should review the CPU usage report produced by CPExpert at the end of the normal rule listing. Compare the CPU time used by the service class identified by Rule WLM222 with the CPU time used by other service classes. Pay particular attention to CPU time used by any service classes at the same or lower importance, to see whether these service classes should receive the CPU service indicated.

- 
- Alternatively, you may wish to increase the Capacity Maximum specified for the resource group. Since applications executing in the service class are being delayed because of CPU capping, you may remove or decrease the delay by increasing the Capacity Maximum for the resource group.
  - Alternatively, you may wish to review the applications executing in the service class identified by Rule WLM222, to determine whether the application code can be optimized so that less CPU time is required.
  - Alternatively, you may wish to examine the CICS region parameters to determine whether appropriate specifications have been provided. For example, the System Initialization Table (SIT) parameters often can significantly alter the amount of CPU time required to support CICS transactions.

If you have licensed the CICS Component of CPExpert, you should execute the CICS Component against the region serving the transactions related to the service class missing its performance goal.

- If none of the above alternatives apply and if Rule WLM222 continually is produced for the service class, you may wish to exclude the service class from CPExpert's analysis. There is little point in having findings produced which cannot be acted upon.<sup>2</sup>

---

<sup>2</sup>Please see Section 2 for information on how to exclude service classes from analysis.